Dynamic logic of evidence-based beliefs

van Benthem, J.F.A.K.; Pacuit, E.J.

Published in:
Studia Logica

DOI:
10.1007/s11225-011-9347-x

J. van Benthem
E. Pacuit

Dynamic Logics of Evidence-Based Beliefs

Dedicated to Professor Ryszard Wójcicki on the occasion of his 80th birthday *

Abstract. This paper adds evidence structure to standard models of belief, in the form of families of sets of worlds. We show how these more fine-grained models support natural actions of "evidence management", ranging from update with external new information to internal rearrangement. We show how this perspective leads to new richer languages for existing neighborhood semantics for modal logic. Our main results are relative completeness theorems for the resulting dynamic logic of evidence.

Keywords: Dynamic Logics of Belief Revision, Neighborhood Models for Modal Logic.

1. Introduction and Motivation

A *rational* belief must be grounded in the evidence available to an agent. However, this relation is delicate, and it raises interesting philosophical and technical issues. Modeling evidence requires richer structures than found in standard epistemic semantics where the accessible worlds aggregate all reliable evidence gathered so far. Even recent more finely-grained plausibility models ordering the epistemic ranges identify too much: belief is indistinguishable from aggregated *best* evidence. At the opposite extreme, one might model evidence syntactically as "formulas received", but this seems overly detailed, and we we lose the intuition that evidence can be semantic in nature, zooming in on some actual world.

In this paper, we explore an intermediate semantic level, viz. that of *neighborhood semantics*, where evidence is recorded as a family of sets of worlds. Neighborhood models have long been a technical tool for studying weak modal logics. But here, we show how they support a notion of evidence

---

* Ryszard has been an inspirational leader in logical research and a major force in the organization of the field for decades. His influence is felt in many circles in many countries. It is an honour to contribute to this special issue of Studia Logica our article to him.

Special issue in honor of Ryszard Wójcicki on the occasion of his 80th birthday
Edited by **J. Czelakowski, W. Dziobiak, and J. Malinowski**
Received January 23, 2011; *Accepted* July 15, 2011

with matching languages for attitudes based on it, as well as an array of natural actions that transform evidence.[1]

Our paper is a pilot study. We develop the basics of neighborhood models for evidence and belief, and the logics that they support. Then we move to the main theme of the paper, showing how these models support natural actions of "evidence management", ranging from external new information to internal rearrangement, which provide a richer picture of evidence than current dynamic logics of information. This dynamic analysis then feeds back into the design of new static neighborhood logics beyond those explored in the literature. All this is just the start of a larger project, and we indicate some further directions at the end.

2. Evidence in Neighbourhood Models

The semantics of evidence that we will use in this paper is based on *neighborhood* models (cf. [4, Chapter 7] and [18, 10], for modern introductions). For convenience, we will stick with *finite models*, though most of our results are easily generalized to infinite settings. Likewise, we will discuss single-agent models only, again, mainly for convenience.

2.1. The basic models

Let W be a set of possible worlds one of which represents the "actual" situation. An agent gathers evidence about this situation from a variety of sources. To simplify things, we assume these sources provide *binary* evidence, i.e., subsets of W, which may, but need not, contain the actual world. In line with our evidence interpretation, we impose some constraints:

- No evidence set is empty (evidence per se is never contradictory),
- The whole universe W is an evidence set (agents know their 'space').

An additional, often-found property is "monotonicity":

If an agent i has evidence X and $X \subseteq Y$, then i has evidence Y.

To us, this is a property of propositions *supported by* evidence, not of evidence itself. We will model this feature differently later on.

[1] Precursors to our semantics are found in the study of belief revision [23, 13, 8], scientific theories [29], topological models for knowledge [17], and sensor-based models of information-driven agency in AI [27].

Definition 2.1 (Evidence Model). An **evidence model** is a tuple $\mathcal{M} = \langle W, E, V \rangle$ with W a non-empty set of worlds, $E \subseteq W \times \wp(W)$ an evidence relation, and $V : \mathsf{At} \to \wp(W)$ a valuation function. A **pointed evidence model** is a pair $\mathcal{M}, w$ with "actual world" w. When E is a constant function, we get a **uniform evidence model** $\mathcal{M} = \langle W, \mathcal{E}, V \rangle, w$ with $\mathcal{E}$ the fixed family of subsets of W related to each state by E. ◁

We write $E(w)$ for the set $\{X \mid wEX\}$. The above two constraints on the evidence function then become:

(Cons) For each state w, $\emptyset \notin E(w)$.
(Triv) For each state w, $W \in E(w)$.

In what follows, we shall mainly work with uniform evidence models. While this may seem very restrictive, the reader will soon see how much relevant structure can be found even at this level.

As stated before, $E(w)$ need not be closed under supersets. Also, even though evidence pieces are non-empty, their combination through the obvious operation of taking *intersections* need not yield consistent evidence: combining disjoint sets will lead to trouble. But importantly, even though an agent may not be able to consistently combine *all* of her evidence, there will be maximal collections that she can safely put together:

Definition 2.2 (Maximal consistent evidence). A family $\mathcal{X}$ of subsets of W has the **finite intersection property** (f.i.p.) if $\bigcap \mathcal{X} \neq \emptyset$. We say $\mathcal{X}$ has the **maximal f.i.p.** if $\mathcal{X}$ has the f.i.p. but no proper extension of $\mathcal{X}$ does. ◁

We will now develop the logic of this framework. Clearly, families of sets give more detail than information states as single sets of worlds.

2.2. A static logic of evidence and belief

We first introduce a basic logic for reasoning about evidence and beliefs.

Language We start with a modal language close to the literature.

Definition 2.3 (Evidence and Belief Language). Let At be a set of atomic propositions. $\mathcal{L}_0$ is the smallest set of formulas generated by the grammar

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid B\varphi \mid \Box\varphi \mid A\varphi$$

where $p \in \mathsf{At}$. Additional propositional connectives $(\wedge, \to, \leftrightarrow)$ are defined as usual, and the existential modality $E\varphi$ is defined as $\neg A \neg \varphi$. ◁

The interpretation of $\Box\varphi$ is "the agent has evidence that implies φ" (the agent has "evidence for" φ") and $B\varphi$ says that "the agents believes that φ". We include the universal modality ($A\varphi$: "φ is true in all states") for convenience. One can also think of this as a form of knowledge.

Having evidence for φ need not imply belief. In order to believe a proposition φ, an agent must consider *all* her evidence for or against φ. To model the latter scenario, we make use of Definition 2.2.[2]

Semantics We now interpret this language on neighborhood models.

Definition 2.4 (Truth). Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model. Truth of a formula $\varphi \in \mathcal{L}_0$ is defined inductively as follows:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$ (for all $p \in \mathsf{At}$)
- $\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ iff $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models \Box\varphi$ iff there is an X with wEX and for all $v \in X$, $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models B\varphi$ iff for each maximal f.i.p. family $\mathcal{X} \subseteq E(w)$ and for all $v \in \bigcap \mathcal{X}$, $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models A\varphi$ iff for all $v \in W$, $\mathcal{M}, v \models \varphi$

The truth set of φ is the set of worlds $[\![\varphi]\!]_{\mathcal{M}} = \{w \mid \mathcal{M}, w \models \varphi\}$. Standard logical notions of **satisfiability** and **validity** are defined as usual. $\triangleleft$

Various extensions to the above modal language make sense. For instance, our notion of belief is cautious, quantifying over *all* maximal f.i.p's. But we might also say that an agent "boldly believes φ" if there is some maximal f.i.p. $\mathcal{X}$ in the current evidence set with $\bigcap \mathcal{X} \subseteq [\![\varphi]\!]$. We will discuss such extensions below.

2.3. Conditional belief and conditional evidence

Our language still lacks some basic features of many logics of belief. Anticipating the evidence dynamics of Section 4, we now introduce *conditional* belief and evidence: $B^{\varphi}\psi$ and $\Box^{\varphi}\psi$ to obtain the language $\mathcal{L}_1$.[3]

---

[2] Analogous ideas occur with conditionals ([12, 37]) and belief revision ([7, 21]).

[3] As usual, we define absolute belief and evidence: $B\varphi := B^{\top}\varphi$ and $\Box\varphi := \Box^{\top}\varphi$.

Conditional evidence The interpretation of $\Box^{\varphi}\psi$ is "the agent has evidence that ψ is true conditional on φ being true". Now, when conditioning on φ one may have evidence X inconsistent with φ. Thus, we cannot simply intersect each piece of evidence with the truth set of φ. We say that $X \subseteq W$ is **consistent** with φ if $X \cap [\![\varphi]\!]_{\mathcal{M}} \neq \emptyset$. Then we define:

- $\mathcal{M}, w \models \Box^{\varphi}\psi$ iff there is an evidence set $X \in E(w)$ which is consistent with φ such that for all worlds $v \in X \cap [\![\varphi]\!]_{\mathcal{M}}$, $\mathcal{M}, v \models \varphi$.

It is easy to see that $\Box^{\varphi}\psi$ is not equivalent to $\Box(\varphi \to \psi)$. No definition with absolute evidence modalities can work, by the results of Section 3.2.

Conditional belief Defining conditional belief ($B^{\varphi}\psi$) involves "relativizing" an evidence model to some formula φ. Some of the agent's current evidence may be inconsistent with φ (i.e., disjoint with $[\![\varphi]\!]_{\mathcal{M}}$). Such inconsistent evidence must be "ignored":

Definition 2.5 (Relativized maximal overlapping evidence). Let $X \subseteq W$. Given a family $\mathcal{X}$ of subsets of W, the relativization $\mathcal{X}^X$ is the set $\{Y \cap X \mid Y \in \mathcal{X}\}$. We say that a family $\mathcal{X}$ has the **finite intersection property relative to** X (X-f.i.p.) if $\bigcap \mathcal{X}^X \neq \emptyset$. $\mathcal{X}$ has the **maximal** X**-f.i.p.** if $\mathcal{X}$ has X-f.i.p. and no proper extension $\mathcal{X}'$ of X has the X-f.i.p. $\lhd$

When X is the truth set of formula φ, we write "maximal φ-f.i.p." for "maximal $[\![\varphi]\!]_{\mathcal{M}}$-f.i.p." and so on. Now we define conditional belief:

- $\mathcal{M}, w \models B^{\varphi}\psi$ iff for each maximal φ-f.i.p. family $\mathcal{X} \subseteq E(w)$, for each world $v \in \bigcap \mathcal{X}^{\varphi}$, $\mathcal{M}, v \models \psi$

While this base language of evidence models looks rich already, it follows familiar patterns. However, there are further natural evidence modalities, and they will come to light through our later analysis of operations that change current evidence. The latter dynamics is in fact the main topic of this paper, but we first explore its static base logic a bit further.

3. Some Logical Theory: Axiomatization and Definability

Axiomatizing valid inference While complete logics for reasoning about evidence are not our main concern, we do note a few facts.

Fact 3.1. *(i) A satisfies all laws of modal* **S5**, *B satisfies all laws of* **KD**, *and $\Box$ satisfies only the principles of the minimal "classical" modal logic: the rule*

of upward monotonicity holds ("from a theorem $\varphi \to \psi$, infer $\Box\varphi \to \Box\psi$".), but conjunction under the modality: $(\Box\varphi \wedge \Box\psi) \to \Box(\varphi \wedge \psi)$ fails. (ii) The following operator connections are valid, but no other implications hold:

$$\begin{tikzcd} & B\varphi \arrow[dr] & \\ A\varphi \arrow[ur] \arrow[dr] & & E\varphi \\ & \Box\varphi \arrow[ur] & \end{tikzcd}$$

Verifying part (i) is straightforward. In particular, the point about $\Box$ is typical for neighborhood semantics: the basic evidence modality does not allow of automatic "aggregation by conjunction". For aggregation to happen, an agent must do *work* – as we will see later on. Part (ii) follows directly using our two basic assumptions (Cons) and (Triv) on evidence.

Over our special class of uniform evidence models, we can say much more. First note that the following are valid:

$$B\varphi \to AB\varphi \quad \text{and} \quad \Box\varphi \to A\Box\varphi.$$

It follows easily that belief introspection is trivially true, as reflected in:

$$\Box\varphi \leftrightarrow B\Box\varphi \quad \text{and} \quad \neg\Box\varphi \leftrightarrow B\neg\Box\varphi$$

These observations suggest the following more general observation:

PROPOSITION 3.2. *On uniform evidence models, each formula of $\mathcal{L}_0$ is equivalent to a formula with modal operator depth 1.*

The proof is essentially as for modal **S5**. Axiomatizing the complete logic of our models seems quite feasible, though the combination of a standard modality B and a neighborhood modality $\Box$ poses some interesting problems. We do not pursue this theme here.

Basic model theory and bisimulation Moving from deductive to expressive power, analyzing definability in our logic requires a notion of bisimulation. Here are a few steps, illustrating our semantics further.

We did not close the set of evidence accepted by an agent at a world under supersets, though our truth definition made the set of propositions that the agent has *evidence for* is closed under weakening. As a concrete illustration, consider the three evidential states pictured below:

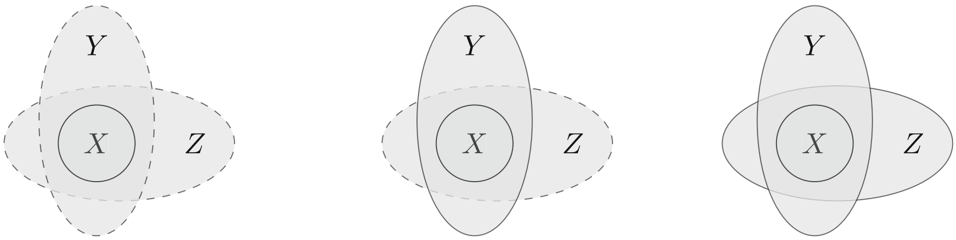

$\mathcal{E}_1 = \{X\}$ $\mathcal{E}_2 = \{X, Y\}$ $\mathcal{E}_3 = \{X, Y, Z\}$

The agent has evidence for X, Y and Z (and all that they entail) in all three cases. Of particular interest is the model $\mathcal{E}_2$, where the agent not only has evidence for Y (because of the accepted evidence X), but also has accepted Y itself as evidence. However, in general, an agent can have *evidence for* X without *accepting* the set X as evidence.

FACT 3.3. *Let the models* $\mathcal{M} = \langle W, E, V \rangle$ *and* $\mathcal{M}' = \langle W, E', V \rangle$ *differ only in their evidence functions. Suppose that, for all* $w \in W$*, (1)* $E(w) \subseteq E'(w)$*, and (2) if* $X' \in E'(w)$*, there is a* $X \in E(w)$ *with* $X \subseteq X'$*. Then, for all* $w \in W$ *and all formulas* $\varphi \in \mathcal{L}_0$*,* $\mathcal{M}', w \models \varphi$ *iff* $\mathcal{M}', w \models \varphi$*.*

The proof is an easy induction on formulas. A natural generalization here is the "monotonic bisimulation" familiar from the literature on neighbourhood semantics [10] and game logics [19].

DEFINITION 3.4 (Monotonic bisimulation). Let $\mathcal{M}_1 = \langle W_1, E_1, V_1 \rangle$ and $\mathcal{M}_2 = \langle W_2, E_2, V_2 \rangle$ be two evidence models. A non-empty relation $Z \subseteq W_1 \times W_2$ is a **bisimulation** if, for all worlds $w_1 \in W_1$ and $w_2 \in W_2$:

Prop If $w_1 Z w_2$, then for all $p \in \mathsf{At}$, $p \in V_1(w_1)$ iff $p \in V_2(w_2)$.

Forth If $w_1 Z w_2$, then for each $X \in E_1^{sup}(w_1)$ there is a $X' \in E_2^{sup}(w_2)$ such that for all $x' \in X'$, there is a $x \in X$ such that xZx'.

Back If $w_1 Z w_2$, then for each $X \in E_2^{sup}(w_2)$ there is a $X' \in E_1^{sup}(w_1)$ such that for all $x' \in X'$, there is a $x \in X$ such that xZx'.

We write $\mathcal{M}_1, w_1 \leftrightarroweq \mathcal{M}_2, w_2$ if there is a bisimulation Z between $\mathcal{M}_1$ and $\mathcal{M}_2$ with $w_1 Z w_2$. A bisimulation Z is **total** if every world in W_1 is related to at least one world in W_2, and vice versa. $\triangleleft$

It is a standard fact that the sublanguage of $\mathcal{L}_0$ without belief modalities is invariant under total bisimulations (totality is needed for the universal modality). Thus, with respect to statements about evidential states, two evidence models are the "same" if they are neighborhood bisimilar. But interestingly, beliefs are *not* invariant under this notion of bisimulation.

Fact 3.5. *The belief modality is not definable with only evidence modalities.*

Proof. Consider the following two evidence models:

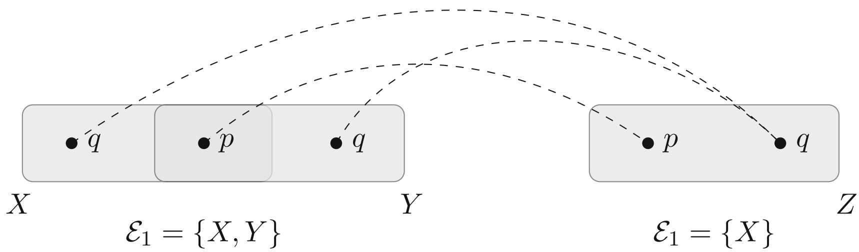

The dashed line is a total bisimulation between the two models. Still, note that Bp is true in the model on the left, but not in that on the right. ■

Finding a notion of bisimulation respecting the whole language of evidence and belief (and their conditionalized variants) is a natural open problem.

4. Evidence Dynamics

Evidence is not a static substance. It is continually affected by new incoming information, and also by processes of internal re-evaluation. Our main point in this paper is to show how this dynamics can be naturally made visible on the neighborhood models that we have introduced.

Our methodology in doing so comes from recent dynamic logics of knowledge update [35, 31] and belief revision [28, 2], which model informational actions driving agency. Formally, these actions change current models, viewed as snapshots of an agent's information and attitudes in some relevant process over time. Examples range from "hard" information change provided by *public announcements* or public observations [20] to softer signals encoding different policies of belief revision (cf. [22]) by radical or more conservative *upgrades* of plausibility orderings. Other dynamic logics describe acts of inference or introspection that raise "awareness" [36, 34], and of questions that modify the focus of a current process of inquiry [32].

While most of these systems work on standard possible worlds models, our neighborhood models of evidence and belief suggest a new scope for these methods in dealing with more finely-structured *evidence dynamics*.[4]

[4]Dynamic neighborhood methods have been used in game scenarios: [5, 38].

Deconstructing public announcement For a start, consider the well-known operation of "public announcement" for a formula φ in a model $\mathcal{M} = \langle W, E, V \rangle$. Defining this is straightforward: remove all $\neg\varphi$-worlds, and intersect the old evidence sets with $[\![\varphi]\!]_{\mathcal{M}}$ when consistently possible. But from the more fine-grained perspective of evidence, the event $!\varphi$ can be naturally "deconstructed" into a combination of three distinct actions:

1. **Evidence addition**: the agent accepts that φ is an "admissible" piece of evidence (perhaps on par with the other available evidence).
2. **Evidence removal**: the agent removes any evidence for $\neg\varphi$.
3. **Evidence modification**: the agent incorporates φ into each piece of evidence gathered so far, making φ the most important piece of evidence.

Our richer evidence models allows us to study these operations individually.

4.1. Public announcements

Definition 4.1 (Public Announcement). Let $\mathcal{M} = \langle W, E, V \rangle$ be an evidence model and φ a formula. The model $\mathcal{M}^{!\varphi} = \langle W^{!\varphi}, E^{!\varphi}, V^{!\varphi} \rangle$ has $W^{!\varphi} = [\![\varphi]\!]_{\mathcal{M}}$, for each $p \in \mathsf{At}$, $V^{!\varphi}(p) = V(p) \cap W^{!\varphi}$, and for all $w \in W$,

$$E^{!\varphi}(w) = \{X \mid \emptyset \neq X = Y \cap [\![\varphi]\!]_{\mathcal{M}} \text{ for some } Y \in E(w)\}. \qquad \triangleleft$$

There is a natural matching dynamic modality $[!\varphi]\psi$ stating that "ψ is true after the public announcement of φ":

$$\text{(PA)} \quad \mathcal{M}, w \models [!\varphi]\psi \text{ iff } \mathcal{M}, w \models \varphi \text{ implies } \mathcal{M}^{!\varphi}, w \models \psi.$$

On evidence models, the standard recursion axioms for public announcement are still valid, yielding dynamic equations for evidence change under hard information. The relevant result expresses a form of "relative completeness":

Theorem 4.2. *The dynamic logic of evidence and belief under public announcement is axiomatized completely over the chosen static base logic, given the usual rulse of Necessitation and Replacement of Provable Equivalents, by (a) the minimal modal logic for the separate dynamic modalities, (b) the following set of recursion axioms:*

(PA1)	$[!\varphi]p$	$\leftrightarrow$	$(\varphi \to p)$	$(p \in \mathsf{At})$
(PA2)	$[!\varphi](\psi \wedge \chi)$	$\leftrightarrow$	$([!\varphi]\psi \wedge [!\varphi]\chi)$	
(PA3)	$[!\varphi]\neg\psi$	$\leftrightarrow$	$(\varphi \to \neg[!\varphi]\psi)$	
(PA4)	$[!\varphi]\Box\psi$	$\leftrightarrow$	$(\varphi \to \Box^{\varphi}[!\varphi]\psi)$	
(PA5)	$[!\varphi]B\psi$	$\leftrightarrow$	$(\varphi \to B^{\varphi}[!\varphi]\psi)$	
(PA6)	$[!\varphi]\Box^{\alpha}\psi$	$\leftrightarrow$	$(\varphi \to \Box^{\varphi\wedge[!\varphi]\alpha}[!\varphi]\psi)$	
(PA7)	$[!\varphi]B^{\alpha}\psi$	$\leftrightarrow$	$(\varphi \to B^{\varphi\wedge[!\varphi]\alpha}[!\varphi]\psi)$	
(PA8)	$[!\varphi]A\psi$	$\leftrightarrow$	$(\varphi \to A[!\varphi]\psi)$	

Table 2: Public Announcement Recursion Axioms

PROOF. We only verify $PA6$ and leave the other axioms to the reader.[5] Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model. Suppose for simplicity that $\mathcal{M}, w \models \varphi$. Then we get

$$
\begin{array}{lll}
\mathcal{M}, w \models [!\varphi]\Box^{\alpha}\varphi & \text{iff} & \mathcal{M}^{!\varphi}, w \models \Box^{\alpha}\varphi \\
& \text{iff} & \text{there is } X \in E^{!\varphi}(w) \text{ compatible with } [\![\alpha]\!]_{\mathcal{M}^{!\varphi}} \\
& & \text{such that } X \cap [\![\alpha]\!]_{\mathcal{M}^{!\varphi}} \subseteq [\![\psi]\!]_{\mathcal{M}^{!\varphi}} \\
& & (\text{note } [\![\psi]\!]_{\mathcal{M}^{!\varphi}} = [\![[!\varphi]\psi]\!]_{\mathcal{M}} \text{ and } [\![\alpha]\!]_{\mathcal{M}^{!\varphi}} = [\![[!\varphi]\alpha]\!]_{\mathcal{M}}) \\
& \text{iff} & \text{there is } X \in E^{!\varphi}(w) \text{ compatible with } [\![[!\varphi]\alpha]\!]_{\mathcal{M}} \\
& & \text{such that } X \cap [\![[!\varphi]\alpha]\!]_{\mathcal{M}} \subseteq [\![[!\varphi]\psi]\!]_{\mathcal{M}} \\
& & (\text{note that } X = Y \cap [\![\varphi]\!]_{\mathcal{M}} \text{ for some } Y \in E(w)) \\
& \text{iff} & \text{there is } Y \in E(w) \text{ compatible with } [\![\varphi \wedge [!\varphi]\alpha]\!]_{\mathcal{M}} \\
& & \text{such that } X \cap [\![\varphi \wedge [!\varphi]\alpha]\!]_{\mathcal{M}} \subseteq [\![[!\varphi]\psi]\!]_{\mathcal{M}} \\
& \text{iff} & \mathcal{M}, w \models \Box^{\varphi\wedge[!\varphi]\alpha}[!\varphi]\psi.
\end{array}
$$

■

4.2. Evidence addition

Next consider the first component in our earlier deconstruction.

DEFINITION 4.3 (Evidence Addition). Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model, and φ a formula in $\mathcal{L}_1$.[6] The model $\mathcal{M}^{+\varphi} = \langle W^{+\varphi}, E^{+\varphi}, V^{+\varphi}\rangle$ has $W^{+\varphi} = W$, $V^{+\varphi} = V$ and for all $w \in W$,

$$E^{+\varphi}(w) = E(w) \cup \{[\![\varphi]\!]_{\mathcal{M}}\}. \qquad \triangleleft$$

[5] These may be validated in analogy with the calculations in [38, Chapter 3].

[6] Eventually, we can even allow formulas from our dynamic evidence logics themselves.

This operation is described by a dynamic modality $[+\varphi]\psi$ stating that "ψ is true after φ is accepted as an admissible piece of evidence":

$$\text{(EA)} \qquad \mathcal{M}, w \models [+\varphi]\psi \text{ iff } \mathcal{M}, w \models E\varphi \text{ implies } \mathcal{M}^{+\varphi}, w \models \psi.$$

Here, since evidence sets are non-empty, the precondition is that φ is true at some state. By contrast, public announcement required that φ be true.

To capture evidence addition in logical terms, we want to find recursion axioms that describe the effect of its action on models ('dynamic equations' of evidence change). Here are a few notions that will be used in our analysis:

Definition 4.4 (Compatible/Incompatible). Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model, $\mathcal{X} \subseteq E(w)$ a family of evidence sets, and φ a formula:

1. $\mathcal{X}$ is maximally φ-**compatible** provided $\cap\mathcal{X} \cap [\![\varphi]\!]_{\mathcal{M}} \neq \emptyset$ and no proper extension $\mathcal{X}'$ of $\mathcal{X}$ has this property; and
2. $\mathcal{X}$ is **incompatible** with φ if there are $X_1, \ldots, X_n \in \mathcal{X}$ such that $X_1 \cap \cdots \cap X_n \subseteq [\![\neg\varphi]\!]_{\mathcal{M}}$. $\lhd$

Maximal $\neg\varphi$-compatibility need not imply incompatibility with φ.

Next, we rephrase our definition of conditional belief, in a new notation:

$$\mathcal{M}, w \models B^{+\varphi}\psi \text{ iff for each maximally } \varphi\text{-compatible } \mathcal{X} \subseteq E(w),$$
$$\bigcap \mathcal{X} \cap [\![\varphi]\!]_{\mathcal{M}} \subseteq [\![\psi]\!]_{\mathcal{M}}$$

But we also need a new conditional belief operator, based on *incompatibility*:

$$\mathcal{M}, w \models B^{-\varphi}\psi \text{ iff for all maximal f.i.p., if } \mathcal{X} \text{ is incompatible with } \varphi \text{ then}$$
$$\bigcap \mathcal{X} \subseteq [\![\psi]\!]_{\mathcal{M}}$$

Now, here is the axiom for belief after evidence addition that we are after:

Lemma 4.5. $[+\varphi]B\psi \leftrightarrow E\varphi \to (B^{+\varphi}[+\varphi]\psi \wedge B^{-\varphi}[+\varphi]\psi)$ *is valid.*

Proof. Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model and φ a formula with $[\![\varphi]\!]_{\mathcal{M}} \neq \emptyset$. We first note the following facts:

1. $\mathcal{X} \subseteq E(w)$ is maximally φ-compatible iff $\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\} \subseteq E^{+\varphi}(w)$ is a maximal f.i.p.
2. $\mathcal{X} \subseteq E(w)$ is a maximal f.i.p. that is incompatible with φ iff $\mathcal{X} \subseteq E^{+\varphi}(w)$ is a maximal f.i.p. that does not contain $[\![\varphi]\!]_{\mathcal{M}}$.

The proof of both facts follows by noting that $E(w) \subseteq E^{+\varphi}(w)$, while any $\mathcal{X}$ that is a maximal f.i.p. in $E^{+\varphi}(w)$ but not in $E(w)$ *must* contain $[\![\varphi]\!]_{\mathcal{M}}$.

Now suppose that $\mathcal{M}, w \models [+\varphi]B\psi$. Then,

$$(*) \qquad \text{for all maximal f.i.p. } \mathcal{X} \subseteq E^{+\varphi}(w), \text{ we have } \bigcap \mathcal{X} \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$$

We must show $\mathcal{M}, w \models B^{+\varphi}[+\varphi]\psi \wedge B^{-\varphi}[+\varphi]\psi$. To see that the left conjunct is true, let $\mathcal{X} \subseteq E(w)$ be any maximally φ-compatible collection of evidence. By (1), $\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\} \subseteq E^{+\varphi}(w)$ is a maximal f.i.p. set. Then, we have

$$\bigcap \mathcal{X} \cap [\![\varphi]\!]_{\mathcal{M}} = \bigcap(\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\}) \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}} = [\![[+\varphi]\psi]\!]_{\mathcal{M}}$$

where the inclusion comes from $(*)$. Since $\mathcal{X}$ was an arbitrary maximally φ-compatible set, we have $\mathcal{M}, w \models B^{+\varphi}[+\varphi]\psi$. For the right conjunct, let $\mathcal{X} \subseteq E(w)$ be any maximal f.i.p. set incompatible with φ. By (2), $\mathcal{X} \subseteq E^{+\varphi}(w)$ is a maximal f.i.p. (not containing $[\![\varphi]\!]_{\mathcal{M}}$). Again by $(*)$,

$$\bigcap \mathcal{X} \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}} = [\![[+\varphi]\psi]\!]_{\mathcal{M}}$$

Hence, since $\mathcal{X}$ was an arbitrary maximal f.i.p. subset of $E(w)$ incompatible with φ, we have $\mathcal{M}, w \models B^{-\varphi}[+\varphi]\psi$. This shows that $[+\varphi]B\psi \to B^{+\varphi}[+\varphi]\psi \wedge B^{+\varphi}[+\varphi]\psi$ is valid.

Suppose now that $\mathcal{M}, w \models B^{+\varphi}[+\varphi]\psi \wedge B^{+\varphi}[+\varphi]\psi$. Then

A. For all maximally φ-compatible $\mathcal{X} \subseteq E(w)$, we have $\bigcap \mathcal{X} \cap [\![\varphi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\psi]\!]_{\mathcal{M}}$; and
B. For all maximally f.i.p. $\mathcal{X} \subseteq E(w)$ incompatible with φ, we have $\bigcap \mathcal{X} \subseteq [\![[+\varphi]\psi]\!]_{\mathcal{M}}$.

We must show $\mathcal{M}^{+\varphi}, w \models B\psi$. Let $\mathcal{X} \subseteq E^{+\varphi}(w)$ be a maximal f.i.p. set. There are two cases to consider. First, $[\![\varphi]\!]_{\mathcal{M}} \in \mathcal{X}$. Then, by (1), $\mathcal{X} - \{[\![\varphi]\!]_{\mathcal{M}}\} \subseteq E(w)$ is maximally φ-compatible. Furthermore, by (A) we have

$$\bigcap \mathcal{X} = \bigcap(\mathcal{X} - \{[\![\varphi]\!]_{\mathcal{M}}\}) \cap [\![\varphi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\psi]\!]_{\mathcal{M}} = [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$$

The second case is $[\![\varphi]\!]_{\mathcal{M}} \notin \mathcal{X}$. Then by (2), $\mathcal{X} \subseteq E(w)$ is a maximal f.i.p. that is incompatible with φ. By (B), we have

$$\bigcap \mathcal{X} \subseteq [\![[+\varphi]\psi]\!]_{\mathcal{M}} = [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$$

In either case, $\bigcap \mathcal{X} \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$; hence, $\mathcal{M}^{+\varphi}, w \models B\psi$, as desired. ∎

This proof will suffice to show that analyzing evidence changes is non-trivial. We had to come up with a new notion of conditional belief.[7]

[7]In particular, the new $B^{-\varphi}\psi$ is not the same as the conditional belief $B^{+\neg\varphi}\psi$.

Language extension But we are not yet done. We have now extended the base language, and hence, we need to find complete recursion axioms for the new conditional beliefs after evidence addition – hopefully, avoiding an infinite regress. To achieve this, let $\mathcal{L}_2$ be the smallest set of formulas generated by the following grammar:

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box\varphi \mid B^{\varphi,\psi}\chi \mid A\varphi$$

where $p \in \mathsf{At}$ and $\overline{\varphi}$ is any finite sequence of formulas from the language.[8]

DEFINITION 4.6 (Truth for $\mathcal{L}_2$). We only define the new modal operator:

$$\mathcal{M}, w \models B^{\varphi,\psi}\chi \text{ iff for all maximally } \varphi\text{-compatible sets } \mathcal{X} \subseteq E(w), \\ \text{if } \bigcap \mathcal{X} \cap [\![\varphi]\!]_{\mathcal{M}} \subseteq [\![\psi]\!]_{\mathcal{M}}, \text{ then } \bigcap \mathcal{X} \cap [\![\varphi]\!]_{\mathcal{M}} \subseteq [\![\chi]\!]_{\mathcal{M}}$$

Note that we can define $B^{+\varphi}$ as $B^{\varphi,\top}$ and $B^{-\varphi}$ as $B^{\top,\neg\varphi}$. ◁

THEOREM 4.7. *The dynamic logic of evidence addition is axiomatized completely by (a) the static base logic of evidence models for the extended language, (b) the minimal modal logic for each separate dynamic modality, and (c) the following set of recursion axioms:*

(EA1)	$[+\varphi]p$	$\leftrightarrow$	$(E\varphi \to p)$ *($p \in \mathsf{At}$)*
(EA2)	$[+\varphi](\psi \wedge \chi)$	$\leftrightarrow$	$([+\varphi]\psi \wedge [+\varphi]\chi)$
(EA3)	$[+\varphi]\neg\psi$	$\leftrightarrow$	$(E\varphi \to \neg[+\varphi]\psi)$
(EA4)	$[+\varphi]\Box\psi$	$\leftrightarrow$	$(E\varphi \to (\Box[+\varphi]\psi \vee A(\varphi \to [+\varphi]\psi)))$
(EA5)	$[+\varphi]B\psi$	$\leftrightarrow$	$(E\varphi \to (B^{+\varphi}[+\varphi]\psi \wedge B^{-\varphi}[+\varphi]\psi))$
(EA6)	$[+\varphi]\Box^{\alpha}\psi$	$\leftrightarrow$	$(E\varphi \to (\Box^{[+\varphi]\alpha}[+\varphi]\psi \vee (E(\varphi \wedge [+\varphi]\alpha) \wedge$ $A((\varphi \wedge [+\varphi]\alpha) \to [+\varphi]\psi))))$
(EA7)	$[+\varphi]B^{\psi,\alpha}\chi$	$\leftrightarrow$	$(E\varphi \to (B^{\varphi\wedge[+\varphi]\psi,[+\varphi]\alpha}[+\varphi]\chi \wedge$ $B^{[+\varphi]\psi,\neg\varphi\wedge[+\varphi]\alpha}[+\varphi]\chi))$
(EA8)	$[+\varphi]A\psi$	$\leftrightarrow$	$(E\varphi \to A[+\varphi]\psi)$

Table 1: Evidence Addition Recursion Axioms

PROOF. For soundness, we explain what the recursion axioms say. The first three axioms express the usual relationship between a dynamic-epistemic

[8]Absolute belief and evidence versions again arise by setting some parameters to $\top$.

modality and boolean connectives. For example, axiom $EA3$ says that evidence addition is *functional*: it maps each evidence model to the unique model representing the situation after the evidence is accepted. Axioms $EA4$ - $EA8$ then describe the precise effect of evidence addition on the agent's (conditional) beliefs and accepted evidence. Axiom $EA4$ says that after accepting φ as evidence, the agent has evidence that ψ just in case either she had evidence for ψ before adding φ, or ψ was implied by φ in the model. Axiom $EA5$ shows the effect of accepting the evidence φ on the agent's beliefs. She comes to believe ψ after accepting φ as evidence just in case she believed ψ conditional on φ being true *and* believes ψ conditional on the incompatibility of φ.[9]

Here are a few more detailed verifications of some key recursion axioms. To simplify the presentation, assume that $[\![\varphi]\!]_{\mathcal{M}} \neq \emptyset$ (so $\mathcal{M}, w \models E\varphi$).

(Axiom $EA4$.) $\mathcal{M}, w \models [+\varphi]\Box\psi$ iff $\mathcal{M}^{+\varphi}, w \models \Box\psi$ iff there is an $X \in E^{+\varphi}(w)$ with $X \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$. By definition, we have $[\![\psi]\!]_{\mathcal{M}^{+\varphi}} = [\![[+\varphi]\psi]\!]_{\mathcal{M}}$. There are two cases to consider for the axiom:

1. $X \in E(w)$. Then, $\mathcal{M}, w \models \Box[+\varphi]\psi$.
2. $X = [\![\varphi]\!]_{\mathcal{M}}$. This means that $[\![\varphi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\psi]\!]_{\mathcal{M}}$ and so $\mathcal{M}, w \models A(\varphi \to [+\varphi]\psi)$.

(Axiom $EA5$.) The validity of this axiom is proven in Lemma 4.5.

(Axiom $EA6$.) $\mathcal{M}, w \models [+\varphi]\Box^{\alpha}\psi$ iff $\mathcal{M}^{+\varphi}, w \models \Box^{\alpha}\psi$ iff there exists $X \in E^{+\varphi}(w)$ consistent with α and having $X \cap [\![\alpha]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$. Again there are two cases:

1. $X \in E(w)$. Then we have $X \cap [\![[+\varphi]\alpha]\!]_{\mathcal{M}} = X \cap [\![\alpha]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}} = [\![[+\varphi]\psi]\!]_{\mathcal{M}}$. Hence $\mathcal{M}, w \models \Box^{[+\varphi]\alpha}[+\varphi]\psi$.
2. $X = [\![\varphi]\!]_{\mathcal{M}}$. Then we have $[\![\varphi]\!]_{\mathcal{M}} \cap [\![\alpha]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$. Therefore, $\mathcal{M}, w \models A((\varphi \wedge [+\varphi]\alpha) \to [+\varphi]\psi)$. Furthermore, since X is consistent with α in $\mathcal{M}^{+\varphi}$, we have $[\![\varphi]\!]_{\mathcal{M}} \cap [\![\alpha]\!]_{\mathcal{M}^{+\varphi}} \neq \emptyset$, and hence $\mathcal{M}, w \models E(\varphi \wedge [+\varphi]\alpha)$.

(Axiom $EA7$.) The proof here is similar to the proof of Lemma 4.5. We first note the following facts:

[9] Contrast this with the much simpler law $PA5$ for *public announcement* which only considers whether the agent believed ψ conditionally on φ.

1. $\mathcal{X} \subseteq E(w)$ is maximally $\varphi \wedge [+\varphi]\psi$-compatible with $\bigcap \mathcal{X} \cap [\![\varphi \wedge [+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\alpha]\!]_{\mathcal{M}}$ iff $\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\} \subseteq E^{+\varphi}(w)$ is maximally ψ-compatible with $\bigcap(\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\}) \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\alpha]\!]_{\mathcal{M}^{+\varphi}}$
2. $\mathcal{X} \subseteq E(w)$ is maximally $[+\varphi]\psi$-compatible with $\bigcap \mathcal{X} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![\neg\varphi \wedge [+\varphi]\alpha]\!]_{\mathcal{M}}$ iff $\mathcal{X} \subseteq E^{+\varphi}(w)$ is maximally ψ-compatible such that $[\![\varphi]\!]_{\mathcal{M}} \notin \mathcal{X}$ and $\bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\alpha]\!]_{\mathcal{M}^{+\varphi}}$.

The proof of these facts is straightforward. For the proof of (2), note that if $\mathcal{X} \subseteq E^{+\varphi}(w)$ is maximally ψ-compatible and $[\![\varphi]\!]_{\mathcal{M}} \notin \mathcal{X}$ then we must have $\bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\neg\varphi]\!]_{\mathcal{M}}$. Suppose that $\mathcal{M}, w \models [+\varphi]B^{\psi,\alpha}\chi$. Then,

$$(*) \qquad \text{for all maximally } \psi\text{-compatible } \mathcal{X} \subseteq E^{+\varphi}(w) \text{ with } \bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\alpha]\!]_{\mathcal{M}^{+\varphi}}, \text{ we have } \bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\chi]\!]_{\mathcal{M}^{+\varphi}}$$

We must show $\mathcal{M}, w \models B^{\varphi\wedge[+\varphi]\psi,[+\varphi]\alpha}[+\varphi]\chi \wedge B^{[+\varphi]\psi,\neg\varphi\wedge[+\varphi]\alpha}[+\varphi]\chi$. For the first conjunct, let $\mathcal{X} \subseteq E(w)$ be maximally $\varphi \wedge [+\varphi]\psi$-compatible with $\bigcap \mathcal{X} \cap ([\![\varphi \wedge [+\varphi]\psi]\!]_{\mathcal{M}}) \subseteq [\![[+\varphi]\alpha]\!]_{\mathcal{M}}$. By (1), $\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\} \subseteq E^{+\varphi}(w)$ is maximally ψ-compatible with $\bigcap(\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\}) \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\alpha]\!]_{\mathcal{M}^{+\varphi}}$. By $(*)$,

$$\bigcap \mathcal{X} \cap [\![\varphi \wedge [+\varphi]\psi]\!]_{\mathcal{M}} = \bigcap \mathcal{X} \cap [\![\varphi]\!]_{\mathcal{M}} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} = \bigcap(\mathcal{X} \cup \{[\![\varphi]\!]_{\mathcal{M}}\}) \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}}$$

$$\subseteq [\![\chi]\!]_{\mathcal{M}^{+\varphi}} = [\![[+\varphi]\chi]\!]_{\mathcal{M}}$$

Hence, $\mathcal{M}, w \models B^{\varphi\wedge[+\varphi]\psi,[+\varphi]\alpha}[+\varphi]\chi$. For the second conjunct, let $\mathcal{X} \subseteq E(w)$ be a maximally $[+\varphi]\psi$-compatible set with

$$\bigcap \mathcal{X} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![\neg\varphi \wedge [+\varphi]\alpha]\!]_{\mathcal{M}}.$$

Then by (2), $\mathcal{X} \subseteq E^{+\varphi}(w)$ is maximally ψ-compatible with $\bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\alpha]\!]_{\mathcal{M}^{+\varphi}}$. By $(*)$, we have

$$\bigcap \mathcal{X} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} = \bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\chi]\!]_{\mathcal{M}^{+\varphi}} = [\![[+\varphi]\chi]\!]_{\mathcal{M}}$$

Hence, $\mathcal{M}, w \models B^{[+\varphi]\psi,\neg\varphi\wedge[+\varphi]\alpha}[+\varphi]\chi$

For the converse, suppose that

$$\mathcal{M}, w \models B^{\varphi\wedge[+\varphi]\psi,[+\varphi]\alpha}[+\varphi]\chi \wedge B^{[+\varphi]\psi,\neg\varphi\wedge[+\varphi]\alpha}[+\varphi]\chi.$$

Then

A. for all maximally $\varphi \wedge [+\varphi]\psi$-compatible sets $\mathcal{X} \subseteq E(w)$ with $\bigcap \mathcal{X} \cap [\![\varphi \wedge [+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\alpha]\!]_{\mathcal{M}}$, we have $\bigcap \mathcal{X} \cap [\![\varphi \wedge [+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\chi]\!]_{\mathcal{M}}$.

B. for all maximally $[+\varphi]\psi$-compatible sets $\mathcal{X} \subseteq E(w)$ with $\bigcap \mathcal{X} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![\neg\varphi \wedge [+\varphi]\alpha]\!]_{\mathcal{M}}$, we have $\bigcap \mathcal{X} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\chi]\!]_{\mathcal{M}}$.

We must show that $\mathcal{M}^{+\varphi}, w \models B^{\psi,\alpha}\chi$. Let $\mathcal{X} \subseteq E^{+\varphi}(w)$ be a maximally ψ-compatible set with $\bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\alpha]\!]_{\mathcal{M}^{+\varphi}}$. There are two cases. First, $[\![\varphi]\!]_{\mathcal{M}} \in \mathcal{X}$. Then by (1), $\mathcal{X} - \{[\![\varphi]\!]_{\mathcal{M}}\}$ is maximally $\varphi \wedge [+\varphi]\psi$-compatible with $\bigcap(\mathcal{X} - \{[\![\varphi]\!]_{\mathcal{M}}\}) \cap [\![\varphi \wedge [+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\alpha]\!]_{\mathcal{M}}$. By (A), we have

$$\bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} = \bigcap(\mathcal{X} - \{[\![\varphi]\!]_{\mathcal{M}}\}) \cap [\![\varphi \wedge [+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\chi]\!]_{\mathcal{M}} = [\![\chi]\!]_{\mathcal{M}^{+\varphi}}$$

The second case is $[\![\varphi]\!]_{\mathcal{M}} \notin \mathcal{X}$. Then, by (2), $\mathcal{X} \subseteq E(w)$ is maximally $[+\varphi]\psi$-compatible with $\bigcap \mathcal{X} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![\neg\varphi \wedge [+\varphi]\alpha]\!]_{\mathcal{M}}$. By (B), we have

$$\bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} = \bigcap \mathcal{X} \cap [\![[+\varphi]\psi]\!]_{\mathcal{M}} \subseteq [\![[+\varphi]\chi]\!]_{\mathcal{M}} = [\![\chi]\!]_{\mathcal{M}^{+\varphi}}$$

In either case, $\bigcap \mathcal{X} \cap [\![\psi]\!]_{\mathcal{M}^{+\varphi}} \subseteq [\![\chi]\!]_{\mathcal{M}^{+\varphi}}$, so $\mathcal{M}^{+\varphi}, w \models B^{\psi,\alpha}\chi$, as desired.

The remainder of the completeness proof follows a standard pattern in dynamic epistemic logic. Working inside out, the stated recursion axioms suffice for successively removing all dynamic modalities from a given formula, leading to a provably equivalent formula in the base language (these steps involve the inference rule of Replacement of Provable Equivalents), whose logic was assumed to be complete. ■

We have now found a complete system of evidence addition with its natural associated static base modalities of conditional belief. This is an interesting extension of standard neighborhood logic by itself. But how does this language fit with our earlier analysis of public announcement? Things turn out to be in harmony.

FACT 4.8. *The following principle suffices for obtaining a complete dynamic logic of evidence addition plus public announcement:*

$$[!\varphi]B^{\psi,\alpha}\chi \leftrightarrow B^{\varphi\wedge[!\varphi]\psi,\varphi\rightarrow[!\varphi]\alpha}[!\varphi]\chi$$

4.3. Evidence removal

Evidence addition and public announcement are two ways in which an agent can incorporate a proposition φ. Public announcement is stronger in that the agent also agrees to *ignore* states inconsistent with φ. The latter attitude is interesting by itself, suggesting an act of *evidence removal*, well-known in studies of belief revision as a natural converse to addition. While modeling "removal" has been a challenge to standard dynamic-epistemic logics, our richer setting suggests a natural logic.

Definition 4.9 (Evidence Removal). Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model, and φ a formula in $\mathcal{L}_1$. The model $\mathcal{M}^{-\varphi} = \langle W^{-\varphi}, E^{-\varphi}, V^{-\varphi}\rangle$ has $W^{-\varphi} = W$, $V^{-\varphi} = V$ and for all $w \in W$,

$$E^{-\varphi}(w) = E(w) - \{X \mid X \subseteq [\![\varphi]\!]_{\mathcal{M}}\}. \qquad \triangleleft$$

This time, the corresponding dynamic modality is $[-\varphi]\psi$ ("after removing the evidence that φ, ψ is true"), defined as follows:

(ER) $\mathcal{M}, w \models [-\varphi]\psi$ iff $\mathcal{M}, w \models \neg A\varphi$ implies $\mathcal{M}^{-\varphi}, w \models \psi$ [10]

Again, we look for a dynamic recursion axiom. As with evidence addition, the analysis is not purely a passive imposition of action superstructure. Finding a total dynamic language that is in harmony again affects the choice of the base language itself, and hence it is an instrument for discovering new logical structure concerning evidence.

For a start, let $\mathcal{L}_1^-$ extend the language $\mathcal{L}_1$ with the operator $[-\varphi]$.

Proposition 4.10. *$\mathcal{L}_1^-$ is strictly more expressive than $\mathcal{L}_1$.*

Proof. Consider evidence models $\mathcal{M}_1 = \langle W, \mathcal{E}_1, V\rangle$ and $\mathcal{M}_2 = \langle W, \mathcal{E}_2, V\rangle$:

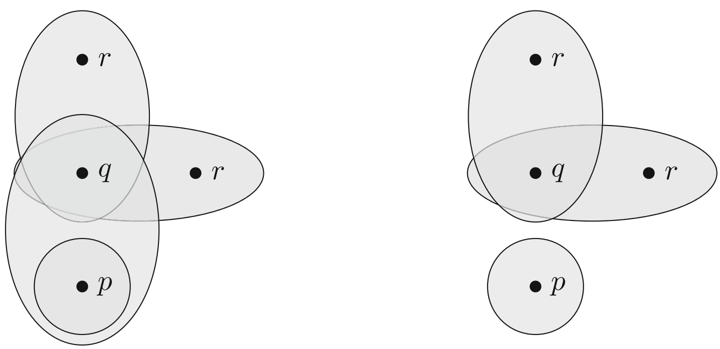

$\mathcal{E}_1$ $\mathcal{E}_2$

The formula $[-p]\Box(p\vee q)$ of $\mathcal{L}_1^-$ is true in $\mathcal{M}_1$ but not in $\mathcal{M}_2$. But no formula of $\mathcal{L}_1$ can distinguish $\mathcal{M}_1$ from $\mathcal{M}_2$. To see this, note that $\mathcal{E}_1^{sup} = \mathcal{E}_2^{sup}$, while the agent has the same beliefs in both models. ■

[10] Removing the evidence for φ is weaker than the usual notion of *contracting* one's beliefs by φ [21]. It is possible to remove the evidence for φ and yet the agent maintains her belief in φ! Formally, $[-\varphi]\neg B\varphi$ is not valid. To see this, let $W = \{w_1, w_2, w_3\}$ with p true only at w_3. Consider an evidence model with two pieces of evidence: $\mathcal{E} = \{\{w_1, w_3\}, \{w_2, w_3\}\}$. The agent believes p and, since the model does not change when removing the evidence for p, $[-p]Bp$ is true. The same is true for the model with explicit evidence for p, i.e., $\mathcal{E}' = \{\{w_1, w_3\}, \{w_2, w_3\}, \{w_3\}\}$.

Adding compatibility So far, we have looked at conditional evidence and beliefs, generalizing the usual notion to restriction and incompatibility versions. This time, we also need to look at evidence that is merely "compatible" with some relevant proposition.

An agent had evidence that ψ conditional on φ if there is evidence consistent with φ such that restriction to the worlds where φ is true entails ψ. Our new conditional operator $\Box_\varphi\psi$ drops the latter condition: it is true if the agent has evidence *compatible with* φ that entails ψ.[11] In general, we include operators $\Box_{\overline{\varphi}}\psi$ where $\overline{\varphi}$ is a sequence of formulas. The evidence entailing ψ must now be compatible with each of $\overline{\varphi}$".

Definition 4.11 (Compatible evidence). Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model and $\overline{\varphi} = (\varphi_1, \ldots, \varphi_n)$ a finite sequence of formulas. A subset $X \subseteq W$ is **compatible with** $\overline{\varphi}$ if, for each φ_i, $X \cap [\![\varphi_i]\!]_\mathcal{M} \neq \emptyset$. $\triangleleft$

Truth of a matching new formula $\Box_{\overline{\varphi}}\psi$ is then defined as follows:

$$\mathcal{M}, w \models \Box_{\overline{\varphi}}\psi \text{ iff some } X \in E(w) \text{ compatible with } \overline{\varphi} \text{ has } X \subseteq [\![\psi]\!]_\mathcal{M}$$

This new operator gives us a very natural reduction axiom for $\Box$: [12]

Fact 4.12. *The formula* $[-\varphi]\Box\psi \leftrightarrow (\neg A\varphi \to \Box_{\neg\varphi}[-\varphi]\psi)$ *is valid.*

Proof. Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model with $[\![\varphi]\!]_\mathcal{M} \neq W$ (otherwise, for all w, $E^{-\varphi}(w) = \emptyset$). We show that $[-\varphi]\Box\psi \leftrightarrow \Box_{\neg\varphi}[-\varphi]\psi$ is valid on $\mathcal{M}$. Let $w \in W$. The key observation is that for all $X \subseteq W$, $X \in E^{-\varphi}(w)$ iff $X \in E(w)$ and X is compatible with $\neg\varphi$. Then we get

$$\begin{array}{lll}
\mathcal{M}, w \models [-\varphi]\Box\varphi & \text{iff} & \mathcal{M}^{-\varphi}, w \models \Box\varphi \\
& \text{iff} & \text{there is a } X \in E^{-\varphi}(w) \text{ such that } X \subseteq [\![\psi]\!]_{\mathcal{M}^{-\varphi}} \\
& & (\text{note that } [\![\psi]\!]_{\mathcal{M}^{-\varphi}} = [\![[-\varphi]\psi]\!]_\mathcal{M}) \\
& \text{iff} & \text{there is a } X \in E(w) \text{ compatible with } \neg\varphi \\
& & \text{such that } X \subseteq [\![[-\varphi]\psi]\!]_\mathcal{M} \\
& \text{iff} & \mathcal{M}, w \models \Box_{\neg\varphi}[-\varphi]\psi.
\end{array}$$

■

But as before, we are not done yet. We also need a reduction axiom for our new operator $\Box_{\overline{\varphi}}$. This can be stated in the same style. But we are not done even then. With the earlier conditional evidence present as well, we need an operator $\Box^\alpha_{\overline{\varphi}}\psi$ saying there is evidence compatible with $\overline{\varphi}$ and α such that the restriction of that evidence to α entails ψ.

---

[11] A set X may be consistent with both φ_1 and φ_2, yet inconsistent with $\varphi_1 \wedge \varphi_2$.

[12] The precondition is needed because the set of all worlds W is an evidence set.

DEFINITION 4.13 (Compatibility evidence - set version). A maximal f.i.p. set $\mathcal{X}$ is **compatible with** a sequence of formulas $\overline{\varphi}$ provided for each $X \in \mathcal{X}$, X is compatible with $\overline{\varphi}$. $\triangleleft$

Language and dynamic logic We are now ready to prceed. Let $\mathcal{L}_3$ be the set of formulas generated by the following grammar:

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid B^{\alpha}_{\overline{\varphi}}\psi \mid \Box^{\alpha}_{\overline{\varphi}}\psi \mid A\varphi$$

where $p \in \mathsf{At}$ and $\overline{\varphi}$ is any finite sequence of formulas from the language.[13]

DEFINITION 4.14 (Truth of $\mathcal{L}_3$). We only define the new modal operators:

- $\mathcal{M}, w \models \Box^{\alpha}_{\overline{\varphi}}\psi$ iff there is a set $X \in E(w)$ compatible with $\overline{\varphi}, \alpha$ such that $X \cap [\![\alpha]\!]_{\mathcal{M}} \subseteq [\![\psi]\!]_{\mathcal{M}}$.
- $\mathcal{M}, w \models B^{\alpha}_{\overline{\varphi}}\psi$ iff for each maximal family α-f.i.p. $\mathcal{X}$ compatible with $\overline{\varphi}$, $\bigcap \mathcal{X}^{\alpha} \subseteq [\![\psi]\!]_{\mathcal{M}}$.

We write $\Box^{\alpha}_{\varphi_1,\ldots,\varphi_n}$ for $\Box^{\alpha}_{(\varphi_1,\ldots,\varphi_n)}$ and $\overline{\varphi}, \alpha$ for $(\varphi_1, \ldots, \varphi_n, \alpha)$. Also, if $\overline{\varphi} = (\varphi_1, \ldots, \varphi_n)$, then we write $[-\varphi]\overline{\varphi}$ for $([-\varphi]\varphi_1, \ldots, [-\varphi]\varphi_n)$. $\triangleleft$

THEOREM 4.15. *The complete dynamic logic of evidence removal is axiomatized, over the complete logic of the static base language for evidence models as enriched above, by the following recursion axioms:*

(ER1)	$[-\varphi]p$	$\leftrightarrow$	$(\neg A\varphi \to p)$	$(p \in \mathsf{At})$
(ER2)	$[-\varphi](\psi \wedge \chi)$	$\leftrightarrow$	$([-\varphi]\psi \wedge [-\varphi]\chi)$	
(ER3)	$[-\varphi]\neg\psi$	$\leftrightarrow$	$(\neg A\varphi \to \neg[-\varphi]\psi)$	
(ER4)	$[-\varphi]\Box^{\alpha}_{\overline{\psi}}\chi$	$\leftrightarrow$	$(\neg A\varphi \to \Box^{[-\varphi]\alpha}_{[-\varphi]\overline{\psi},\neg\varphi}[-\varphi]\chi)$	
(ER5)	$[-\varphi]B^{\alpha}_{\overline{\psi}}\chi$	$\leftrightarrow$	$(\neg A\varphi \to B^{[-\varphi]\alpha}_{[-\varphi]\overline{\psi},\neg\varphi}[-\varphi]\chi)$	
(ER6)	$[-\varphi]A\psi$	$\leftrightarrow$	$(\neg A\varphi \to A[-\varphi]\psi)$	

Table 3: Evidence Removal Recursion Axioms

PROOF. We only do axiom $ER5$. Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model, $w \in W$ and $[\![\varphi]\!]_{\mathcal{M}} \neq W$. First of all, note that the key observation in the

[13] Absolute belief and evidence versions again arise by setting some parameters to $\top$.

proof of Fact 4.12 extends to sets of evidence sets (cf. Definition 4.13). That is, for all worlds w, $\mathcal{X} \subseteq E^{-\varphi}(w)$ is compatible with $\overline{\psi}$ iff $\mathcal{X} \subseteq E(w)$ is compatible with $[-\varphi]\overline{\psi}, \neg\varphi$. Furthermore, for all states w, $\mathcal{X} \subseteq E^{-\varphi}(w)$ is a maximal α-f.i.p. iff $\mathcal{X} \subseteq E(w)$ is a maximal $[-\varphi]\alpha$-f.i.p. *compatible with* $\neg\varphi$.[14] Then we can calculate as follows:

$$\begin{array}{lll} \mathcal{M}, w \models [-\varphi]B^{\alpha}_{\overline{\psi}}\chi & \text{iff} & \mathcal{M}^{-\varphi}, w \models B^{\alpha}_{\overline{\psi}}\chi \\ & \text{iff} & \text{for each maximal } \alpha\text{-f.i.p. } \mathcal{X} \subseteq E^{-\varphi}(w) \\ & & \text{compatible with } \overline{\varphi}, \bigcap \mathcal{X}^{\alpha} \subseteq [\![\chi]\!]_{\mathcal{M}^{-\varphi}} = [\![[-\varphi]\chi]\!]_{\mathcal{M}} \\ & \text{iff} & \text{for each maximal } [-\varphi]\alpha\text{-f.i.p. } \mathcal{X} \subseteq E(w) \\ & & \text{compatible with } [-\varphi]\overline{\varphi} \text{ and } \neg\varphi, \\ & & \bigcap \mathcal{X}^{[-\varphi]\alpha} \subseteq [\![[-\varphi]\chi]\!]_{\mathcal{M}} \\ & \text{iff} & \mathcal{M}, w \models B^{[-\varphi]\alpha}_{\overline{\psi}, \neg\varphi}[-\varphi]\chi. \end{array}$$

■

The above principles state the essence of evidence removal, as well as the beliefs one can still have after such an event. The additional insight is that removal essentially involves compatibility as well as implication between propositions – something of independent logical interest.

Logics for evidence once more This is a beginning rather than an end. Extending the base language in this manner will have repercussions for our earlier analyses. Still, it is possible to also find reduction axioms for our new evidence and belief operators under actions of evidence addition and public announcement. For example, for the compatible evidence operator $\Box_{\overline{\psi}}$ with $\overline{\psi} = (\psi_1, \ldots, \psi_n)$, we have the following validities:

$$[+\varphi]\Box_{\overline{\psi}}\chi \leftrightarrow [E\varphi \to (\Box_{[+\varphi]\overline{\psi}}[+\varphi]\chi \vee (\bigwedge_{i=1,\ldots,n} E(\varphi \wedge \psi_i) \wedge A(\varphi \to [+\varphi]\psi)))]$$

$$[!\varphi]\Box_{\overline{\psi}}\chi \leftrightarrow (\varphi \to \Box^{\varphi}_{[!\varphi]\overline{\psi}}[!\varphi]\chi)$$

We do not include all combinations here. The key point is that the analysis is in harmony not leading to further extensions of the base language.

Perhaps more challenging open problems have to do with the "action algebra" of combining our three basic actions on evidence. What happens

[14]The last clause about being compatible with $\neg\varphi$ is crucial: not every $\mathcal{X} \subseteq E^{-\varphi}(w)$ that is a maximal α-f.i.p. corresponds to a maximal $[-\varphi]\alpha$-f.i.p. subset of $E(w)$.

when we compose them? Our guess is that we need to move to an "event model" version of our logics in the style of dynamic-epistemic logic.

4.4. Evidence modification

We have analyzed the two major operations on evidence that we can see. Nevertheless, the space of potential operations on neighborhood models is much larger, even if we impose conditions of bisimulation invariance as in [10]). Instead of exploring this wide realm, we show one new operation. So far, we added or removed evidence. But one could also *modify* the existing pieces of evidence. To see the difference, here is a new way of making some proposition φ highly important:

DEFINITION 4.16 (Evidence Upgrade). [15] Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model and φ a formula in L_1. The model $\mathcal{M}^{\Uparrow\varphi} = \langle W^{\Uparrow\varphi}, E^{\Uparrow\varphi}, V^{\Uparrow\varphi}\rangle$ has $W^{\Uparrow\varphi} = W$, $V^{\Uparrow\varphi} = V$, and for all $w \in W$,

$$E^{\Uparrow\varphi}(w) = \{X \cup [\![\varphi]\!]_{\mathcal{M}} \mid X \in E(w)\} \cup [\![\varphi]\!]_{\mathcal{M}}. \qquad \triangleleft$$

This is stronger than simply adding $[\![\varphi]\!]_{\mathcal{M}}$ as evidence, since one modifies each admissible evidence set. But it is weaker than publicly announcing φ, as the agent retains the ability to consistently condition on $\neg\varphi$.

FACT 4.17. *The following recursion principles are valid:*

1. $[\Uparrow\varphi]\Box\psi \leftrightarrow (E\varphi \to A(\varphi \to [\Uparrow\varphi]\psi))$
2. $[\Uparrow\varphi]B\psi \leftrightarrow (E\varphi \to A(\varphi \to [\Uparrow\varphi]\psi))$

PROOF. For the second law, note that in $E^{\Uparrow\varphi}(w)$, there is only one maximal f.i.p. whose intersection must be $[\![\varphi]\!]_{\mathcal{M}}$. The proof of the first law goes like that of Fact 4.18 below. ■

As these principles show, $\Uparrow\varphi$ gives a very special status to the incoming information φ, blurring the distinction between evidence and belief. This suggests a weaker operation that modifies the evidence sets in favor of φ, but does not add explicit support for φ. Define $\mathcal{M}^{\Uparrow_w\varphi}$ as in Definition 4.16 except for setting $E^{\Uparrow_w\varphi}(w) = \{X \cup [\![\varphi]\!]_{\mathcal{M}} \mid X \in E(w)\}$. A simple modification to Principle 2 in the above fact gives us a valid principle for our evidence operator. However, the case of belief poses some problems.[16]

---

[15]This operation is a bit like "radical upgrade" in dynamic logics of belief change.

[16]The new complication is that, without adding φ to the evidence sets, intersections of maximal f.i.p. sets in the upgraded model may contain more than just φ states.

FACT 4.18. *The formula* $[\Uparrow_w\varphi]\Box\psi \leftrightarrow (\Box[\Uparrow_w\varphi]\psi \wedge A(\varphi \to [\Uparrow_w\varphi]\psi))$ *is valid.*

PROOF. Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model with $w \in W$. Then,

$$
\begin{array}{lll}
\mathcal{M}, w \models [\Uparrow_w\varphi]\Box\psi & \text{iff} & \mathcal{M}^{\Uparrow\varphi}, w \models \Box\psi \\
& \text{iff} & \text{there is a } X \in E^{\Uparrow\varphi}(w) \text{ such that } X \subseteq [\![\psi]\!]_{\mathcal{M}^{-\varphi}} \\
& & (\text{note that } [\![\psi]\!]_{\mathcal{M}^{\Uparrow\varphi}} = [\![[\Uparrow\varphi]\psi]\!]_{\mathcal{M}}) \\
& \text{iff} & \text{there is } X' \in E(w) \text{ with} \\
& & X' \cup [\![\varphi]\!]_{\mathcal{M}} = X \subseteq [\![[\Uparrow\varphi]\psi]\!]_{\mathcal{M}} \\
& \text{iff} & \text{there is } X' \in E(w) \text{ with} \\
& & X' \subseteq [\![[\Uparrow\varphi]\psi]\!]_{\mathcal{M}} \text{ and } [\![\varphi]\!]_{\mathcal{M}} \subseteq [\![[\Uparrow\varphi]\psi]\!]_{\mathcal{M}} \\
& \text{iff} & \mathcal{M}, w \models \Box[\Uparrow\varphi]\psi \wedge A(\varphi \to [\Uparrow\varphi]\psi)
\end{array}
$$

■

4.5. From external to internal actions: evidence combination

We have now brought to light a rich repertoire of evidence-modifying actions. Still, the operations discussed above all exemplify "external evidence dynamics" responding to some outside source, where the agent reacts appropriately, either by incorporating φ or removing φ from consideration. But our neighborhood models also suggest *internal operations* that arise from pondering the evidence, without external triggers. We will discuss only one such internal operation in this paper, be it a basic one.

One natural operation available to an agent is to *combine* her evidence. Of course, as we have noted, an agent's evidence may be contradictory, so she can only combine evidence that is not inconsistent.

DEFINITION 4.19 (Evidence combination). Let $\mathcal{M} = \langle W, E, V\rangle$ be an evidence model. The model $\mathcal{M}^{\#} = \langle W^{\#}, E^{\#}, V^{\#}\rangle$ has $W^{\#} = W$, $V^{\#} = V$ and for all $w \in W$, $E^{\#}(w)$ is the smallest set closed under (non-empty) intersection and containing $E(w)$. The corresponding modal operator is defined as $\mathcal{M}, w \models [\#]\varphi$ iff $\mathcal{M}^{\#}, w \models \varphi$. ◁

A complete study of this operation will be left for future work, since it poses some challenges to our recursive style of analysis so far.[17] Nevertheless, we can observe the following interesting facts:

FACT 4.20. *The following formulas are valid on evidence models:*

[17]The problem may be that standard modal languages are too poor, forcing us upward in expressive power to hybrid or first-order logics – but we suspend judgment here.

1. *$\Box[\#]\varphi \to [\#]\Box\varphi$ (combining evidence does not remove old evidence[18])*
2. *$B[\#]\varphi \leftrightarrow [\#]B\varphi$ (beliefs are immune to evidence combination)*
3. *$B\varphi \to [\#]\Box\varphi$ (beliefs are explicitly supported after combining evidence[19]*
4. *For factual φ, $B\varphi \to \neg[\#]\Box\neg\varphi$ (if an agent believes φ then the agent cannot combine her evidence so that there is evidence for $\neg\varphi$)*

Proof. The proof that the first three items are valid is left to the reader. For the fourth, note that $\Box\neg\varphi \to \neg B\varphi$ is valid. The proof is as follows: First of all, in any evidence model $\mathcal{M} = \langle W, E, V\rangle$, every piece of evidence in $X \in E(w)$ is contained in a maximal f.i.p. $\mathcal{X} \subseteq E(w)$ (models are finite, so simply find the maximal f.i.p. containing X which may be $\{X, W\}$). Suppose that $\Box\neg\varphi$ is true at a state w, then there is an $X \in E(w)$ such that $X \subseteq [\![\neg\varphi]\!]_{\mathcal{M}}$. Let $\mathcal{X}$ be the maximal f.i.p. containing X. Hence, $\bigcap\mathcal{X} \subseteq X \subseteq [\![\neg\varphi]\!]_{\mathcal{M}}$. Therefore, $B\varphi$ is not true at w. This shows that $\Box\neg\varphi \to \neg B\varphi$ is valid, as desired. We can then derive Principle 3 by noting the following series of implications:

$$B\varphi \to [\#]B\varphi \to [\#]\neg\Box\neg\varphi \to \neg[\#]\Box\neg\varphi$$

Here the first implication follows from the second principle applied to factual formulas φ (for which $\varphi \leftrightarrow [\#]\varphi$ is valid), the second implication follows from the fact that $B\varphi \to \neg\Box\neg\varphi$ is valid (as argued above) while $[\#]$ is a normal modal operator and the third implication follows from the fact that the evidence combination operation is functional. ■

5. Comparison with plausibility models

In this section, we will contrast our neighborhood models with another modal framework for belief change. This excursion (which can be skipped without loss of coherence) throws new light on our earlier proposals. We merely state some main notions and results without proof, referring to [33] for details.

---

[18] Definition 4.19 assumed that, for all states w, $E(w) \subseteq E^{\#}(w)$. Thus, in the process of combination, an agent does not notice any inconsistencies that may be present in her evidential state. A deeper analysis would include acts of removing observed inconsistencies.

[19] The converse is not valid. In fact, one can read the combination $[\#]\Box$ as an existential version of our belief operator. It is true if there is some maximal collection of evidence whose intersection implies φ. In plausibility models for doxastic logic, this says that φ is true throughout *some maximal cluster*. This notion of belief is much riskier then $B\varphi$, and again we encounter the variety of agent attitudes mentioned in Section 2.

Plausibility models Originally used as a semantics for conditionals (cf. [14]), the following idea is wide-spread in modal logics of belief [28, 31, 2, 8]. One endows epistemic ranges with an ordering $w \preceq v$ of relative *plausibility* on worlds (usually uniform across indistinguishable worlds): "according to the agent, world v is at least as plausible as w".[20] Plausibility orders are usually reflexive and transitive, and often also *connected*, making every two worlds comparable. Our discussion will allow pre-orders with incomparable worlds. Plausibility models support a general dynamics of informational action through *model change*, for which we refer to cf. [31].

DEFINITION 5.1 (Plausibility model). A **plausibility model** is a tuple $\mathcal{M} = \langle W, \preceq, V \rangle$ where W is a finite nonempty set, $\preceq \subseteq W \times W$ is a reflexive and transitive ordering on W, and $V : \mathsf{At} \to \wp(W)$ is a valuation function. If $\preceq$ is also *connected* (for each $w, v \in W$, either $w \preceq v$ or $v \preceq w$) then we say $\mathcal{M}$ is a **connected plausibility model**. A pair $\mathcal{M}, w$ where w is a state is called a **pointed (connected) plausibility model**. $\triangleleft$

Language and logic Plausibility models interpret a standard doxastic language. Let $\mathcal{L}_{\preceq}$ be the smallest set of formulas generated by the syntax

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid B^{\varphi}\psi \mid [\preceq]\varphi \mid A\varphi$$

To define truth, we need some notation. For $X \subseteq W$, let

$$Min_{\preceq}(X) = \{v \in X \mid v \preceq w \text{ for all } w \in X\}$$

Given a set X, $Min_{\preceq}(X)$ is the set of most plausible worlds in X (i.e., minimal elements of X according to the plausibility order). We only consider the modal operators:

- $\mathcal{M}, w \models B^{\varphi}\psi$ iff $Min_{\preceq}([\![\varphi]\!]_{\mathcal{M}}) \subseteq [\![\psi]\!]_{\mathcal{M}}$
- $\mathcal{M}, w \models [\preceq]\varphi$ iff for all $v \in W$, if $v \preceq w$ then $\mathcal{M}, v \models \varphi$
- $\mathcal{M}, w \models A\varphi$ iff for all $v \in W$, $\mathcal{M}, v \models \varphi$.

With $B\varphi$ defined as $B^{\top}\varphi$, we get the usual notion of belief as truth in all minimal worlds. We can think of this as follows. Any pre-order forms a partial order of "clusters", maximal subsets where the relation is universal. A finite pre-order has one or more final clusters, not having any proper successors. (If the order is connected, there is only one final cluster.) Belief means truth in all final clusters.

---

[20] In conditional semantics, plausibility or "similarity" is a world-dependent order.

The logic of this system is basically the minimal conditional logic over pre-orders that we have seen before. Next, consider definability. Plausibility orders are binary relations supporting a standard modal language. Indeed, [3] showed how belief and conditional belief are definable in the language with A and $[\preceq]$ only:

FACT 5.2. *Belief and conditional belief can be explicitly defined as follows:*

- $B\varphi \;:=\; A\langle\preceq\rangle[\preceq]\varphi$
- $B^{\varphi}\psi \;:=\; A(\varphi \to \langle\preceq\rangle(\varphi \wedge [\preceq](\varphi \to \psi)))$

While the plausibility modality by itself may look like a technical device, [2] interpret $[\preceq]\varphi$ as "a *safe belief* in φ". Following [26], this amounts to the beliefs an agent retains under all new true information about the actual world.[21] Indeed, this simple modal language over plausibility models will turn out to be a natural stage of expressive power.

From plausibility models to evidence models Here is an intuitive connection. Let $\mathcal{M} = \langle W, \preceq, V\rangle$ be a plausibility model:

the appropriate evidence sets are the downward $\preceq$-closed sets of worlds.

To be more precise, we fix some notation:

- Given a $X \subseteq W$, let $X{\downarrow_{\preceq}} = \{v \in W \mid \exists x \in X \text{ and } v \preceq x\}$ (we write $X{\downarrow}$ when it is clear which plausibility ordering is being used).
- A set $X \subseteq W$ is $\preceq$-closed if $X{\downarrow_{\preceq}} \subseteq X$.

Here is the formal definition for the above simple idea:

DEFINITION 5.3 (Plausibility-Based Evidence Model). Let $\mathcal{M} = \langle W, \preceq, V\rangle$ be a plausibility model. The **evidence model generated from $\mathcal{M}$** is[22] $EV(\mathcal{M}) = \langle W, \mathcal{E}_{\preceq}, V\rangle$ with $\mathcal{E}_{\preceq}$ as follows:

$$\mathcal{E}_{\preceq} = \{X \mid \emptyset \neq X \text{ is } \preceq\text{-closed }\} \qquad \triangleleft$$

Given a plausibility model $\mathcal{M}$, the evidence model generated by $\mathcal{M}$ clearly satisfies the basic properties required in Section 2: the sets are non-empty, and the whole universe is among them. But more can be said:

FACT 5.4. *Evidence sets of models $EV(\mathcal{M})$ are closed under intersections.*

[21] For the same notion in the computational literature on agency, cf. [25].

[22] Here the set of worlds and valuation function remain as in the model $\mathcal{M}$.

Plausibility models represent a situation where the agent has already "combined" all of her evidence (cf. the operation # studied in Section 3.4), as reflected in this technical property:

$$\text{If } X, Y \in \mathcal{E}_{\preceq} \text{ and } X \cap Y \neq \emptyset \text{ then } X \cap Y \in \mathcal{E}_{\preceq}.$$

This connection between plausibility models and evidence models can be extended to a translation between their languages:

Definition 5.5 (P-translation). The translation $(\cdot)^P : \mathcal{L} \to \mathcal{L}_{\preceq}$ is given by:

- $p^P = p$, $(\neg\varphi)^P = \neg\varphi^P$, $(\varphi \wedge \psi)^P = \varphi^P \wedge \psi^P$, $(A\varphi)^P = A\varphi^P$,
- $(\Box\varphi)^P = E[\preceq]\varphi^P$,
- $(\Box^{\varphi}\psi)^P = E\langle\preceq\rangle(\varphi^P \wedge [\preceq](\varphi^P \to \psi^P))$,
- $(\Box^{\varphi}_{\vec{\gamma}}\psi)^P = E(\bigwedge_i \langle\preceq\rangle\gamma_i^P \wedge \langle\preceq\rangle(\varphi^P \wedge [\preceq](\varphi^P \to \psi^P)))$,
- $(B^{\varphi}\psi)^P = A(\varphi^P \to (\langle\preceq\rangle\varphi^P \wedge [\preceq](\varphi^P \to \psi^P)))$,
- $(B^{\varphi,\alpha}\psi)^P = A(([\preceq]\alpha^P \wedge [\preceq]\langle\preceq\rangle\varphi^P) \to \langle\preceq\rangle(\varphi^P \wedge [\preceq](\varphi^P \to \psi^P)))$, and
- $(B^{\varphi}_{\vec{\gamma}}\psi)^P = A((\varphi^P \wedge \bigwedge_i\langle\preceq\rangle\gamma_i^P) \to (\langle\preceq\rangle(\varphi^P \wedge \bigwedge_i\langle\preceq\rangle\gamma_i^P) \wedge [\preceq]((\varphi^P \wedge \bigwedge_i\langle\preceq\rangle\gamma_i^P) \to \psi^P)))$. ◁

Lemma 5.6. *Let $\mathcal{M} = \langle W, \preceq, V\rangle$ be a plausibility model. For any formula $\varphi \in \mathcal{L}_1$ and world $w \in W$,*

$$\mathcal{M}, w \models \varphi^P \text{ iff } EV(\mathcal{M}), w \models \varphi$$

Details of this connection can be found in [33].

From evidence models to plausibility models Going in the opposite direction, given a family of evidence sets, how to induce a plausibility order? We use a notion from topology, and models of relation merge (cf. [1, 15]), the "specialization (pre)-order":

Definition 5.7 (Plausibility based evidence model). Suppose that $\mathcal{M} = \langle W, \mathcal{E}, V\rangle$ is an (uniform) evidence model. The plausibility model generated by $\mathcal{M}$ is the structure $ORD(\mathcal{M}) = \langle W, \preceq_E, V\rangle$ where $\preceq_E$ is an ordering on W defined as follows:[23]

$$w \preceq_{\mathcal{E}} v \text{ iff } \forall X \in \mathcal{E},\ v \in X \text{ implies } w \in X$$ ◁

[23] $\preceq_E$ is reflexive and transitive, so $ORD(\mathcal{M})$ is indeed a plausibility model.

Our two representations are related as follows:

FACT 5.8. *(i) For all models plausibility models $\mathcal{M}$, $ORD(EV((\mathcal{M})) = \mathcal{M}$, (ii) The identity $EV(ORD(\mathcal{M})) = \mathcal{M}$ does not hold for all evidence models $\mathcal{M}$. (iii) For all evidence models $\mathcal{M}$, $EV(ORD(\mathcal{M})) = \mathcal{M}^{\#}$, where # is the combination operation of Section 4.19.*

Translations and languages The preceding connection again matches a translation between modal languages, in particular for (conditional) beliefs on evidence models and their induced plausibility models. But other notions are less easily reduced. For instance, [33] show how dealing with safe belief on plausibility orders requires a new notion of *reliable evidence* that extends our earlier evidence logics.

This concludes our brief comparison of relational and neighborhood semantics for belief and evidence. We have clarified their relationship as one of generalization, where neighborhood models describe one more level of detail: the internal combination stages for evidence. Even so, many of the new operations that we have found in earlier sections would also make sense as definable operators in the natural modal logic of plausibility models, and we have shown how various interesting new questions arise at this interface.

6. Conclusion and Further Directions

We have shown that evidence dynamics on neighborhood models offers a rich environment for modeling information flow and pursuing logical studies. Here are some avenues for further research. Some are more technical, some increase coverage. We start with the former.

Exploring the static logic We have found quite a few new evidence-based modalities of conditional belief. What is the complete logic of this system? This is a new question of axiomatization, that can be appreciated also outside of our dynamic perspective. One reason for its complexity may be that we are mixing a language of neighborhood-based modalities with normal operators of belief with a matching relational semantics.

Notions of bisimulation Our languages invite matching notions of structural invariance for evidence models. We already saw that standard bisimulation for neighborhood models matches modal logics with only evidence operators. But Fact 3.5 showed that this does not extend to modalities of

belief referring to intersections of maximally consistent families of evidence sets. And we introduced even stronger modal languages in our discussion of dynamics in Section 4. What stronger notions of bisimulation respect this evidence structure?

Finally, there are also obvious technical generalizations to be made: to *infinite models*, and to DEL-style *product update* mechanisms for rich input.

Reliable evidence and its sources But our setting can also model further phenomena. For instance, there is a natural notion of "reliable" evidence, based only on sets containing the actual world. What is the complete logic of this operator? This suggests a broader study of belief based on reliable evidence, in line with current trends in epistemology. But eventually, we also want to introduce explicit modeling of *sources of evidence* and what agents know or believe about their reliability.

Social notions We have seen how rich evidence structure arises in single agent models. But *multi-agent scenarios* are also natural: e.g., argumentation is social confrontation of evidence, which may result in new group attitudes among participants. This raises interesting issues about natural notions of *group evidence* and belief. Here evidence structure soon takes us beyond the usual notions from the epistemic literature based on relational models.

To illustrate this, suppose there are two agents i and j and a multi-agent uniform evidence model $\mathcal{M} = \langle W, \mathcal{E}_i, \mathcal{E}_j, V \rangle$. We can then ask what evidence the group i, j has. One option here is mere throwing together into one "unprocessed" new evidence set:

$$\mathcal{M}, w \models \Box^{\{i,j\}}\varphi \text{ iff there is a } X \in \mathcal{E}_i \cup \mathcal{E}_j \text{ such that } X \subseteq [\![\varphi]\!]_{\mathcal{M}}$$

At another extreme, the group might only take into account only evidence that is *shared* between the two agents:

$$\mathcal{M}, w \models \Box_{\{i,j\}}\varphi \text{ iff there is a } X \in \mathcal{E}_i \cap \mathcal{E}_j \text{ such that } X \subseteq [\![\varphi]\!]_{\mathcal{M}}$$

For a rough analogy on relational models, the latter form of group evidence is more like common knowledge, while the former is more related to distributed knowledge. But our richer neighborhood models also allow for further distinctions. In particular, the agents can also "pool" their evidence creating a *new* evidential states by *combining* their evidence:

$$\mathcal{E}_i \sqcap \mathcal{E}_j = \{Y \mid \emptyset \neq Y = X \cap X' \text{ with } X \in \mathcal{E}_i \text{ and } X' \in \mathcal{E}_j \ \}$$

We can then define group evidence and belief modalities for this evidential state as we did in Section 2.2. For instance,

$$\mathcal{M}, w \models [i \sqcap j]\varphi \text{ iff there exists } X \in \mathcal{E}_i \sqcap \mathcal{E}_j \text{ with } X \subseteq [\![\varphi]\!]_{\mathcal{M}}$$

What other natural notions of group evidence can we find?

Priority structures The evidence dynamics in this paper treats evidence sets on a par. As a consequence, removal may seem arbitrary and non-deterministic, since there is nothing in the structure of the evidence itself which directs the process. A next reasonable step would be to model *levels of reliability* of evidence. One natural format for this are the "priority graphs" of [1], which have already been used extensively in dynamic-epistemic logic [15, 8]. These graphs provide much richer input to evidence management, and can break stalemates between conflicting pieces of evidence. It should be possible to extend the above framework to one with ordered evidence sets – and conversely, then, our logics may help provide something that has been missing so far: modal logics working directly on priority graphs.

Other logics of evidence "Evidence" is a notion with many different aspects. Our proposal has been set-theoretic and semantic, while there are many other treatments of evidence for a proposition φ, in terms of proofs for φ, or using the balance of probability for φ versus $\neg\varphi$. What we find particularly pressing is a junction with more syntactic approaches making evidence something coded that can be operated on in terms of inference and computation. If finer operational aspects of inference and introspection enter one's notion of evidence, then the methods of this paper should be extended to include dynamic logics of awareness and inference [30, 36].

Related frameworks Finally, the analysis in this paper should be linked up with other traditions, including the seminal work by [6] and [24] on evidence, probabilistic logics of evidence [9], or the "topologic" of [17]. And one can add the "priority graphs" inducing preference orders in [16], or the "belief base" account of belief revision in [11]. We intend to clarify these connections in future work.

Conclusion We have made a proposal for using neighborhood models as fine-grained evidence structures that allow for richer representation of information than current relational models of belief. We have shown how these

structures support a rich dynamics of evidence change that goes beyond current logics of belief revision. A number of relative completeness theorems identified the key dynamic equations governing this process, while also suggesting new static languages of evidence and belief. Finally, we discussed some new issues that lie ahead, such as adding priority structure and group evidence, exploiting the richer neighborhood setting.

References

[1] Andreka, Hajnal, Mark Ryan, and Pierre Yves Schobbens, 'Operators and laws for combining preference relations', *Journal of Logic and Computation* 12(1):13–53, 2002.

[2] Baltag, A., and S. Smets, 'Conditional doxastic models: A qualitative approach to dynamic belief revision', in G. Mints, and R. de Queiroz, (eds.), *Proceedings of WOLLIC 2006, LNCS*, vol. 165, 2006, pp. 5–21.

[3] Boutilier, Craig, *Conditional Logics for Default Reasoning and Belief Revision*, Ph.D. thesis, University of Toronto, 1992.

[4] Chellas, Brian, *Modal Logic: An Introduction*, Cambridge University Press, 1980.

[5] Demey, Lorenz, *Agreeing to Disagree in Probabilistic Dynamic Epistemic Logic.*, Master's thesis, ILLC University of Amsterdam (LDC 2010-14), 2010.

[6] Dempster, A. P., 'Upper and lower probabilities induced by a multivalued mapping', *Annals of Mathematical Statistics* 38(2):325–339, 1967.

[7] Gärdenfors, Peter, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press, 1988.

[8] Girard, Patrick, *Modal Logic for Belief and Preference Change*, Ph.D. thesis, ILLC University of Amsterdam Dissertation Series DS-2008-04, 2008.

[9] Halpern, Joe, and Ricardo Pucella, 'A logic for reasoning about evidence', *Journal of AI Research* 26(1):1–34, 2006.

[10] Hansen, Helle Hvid, Clemens Kupke, and Eric Pacuit, 'Neighbourhood structures: Bisimilarity and basic model theory', *Logical Methods in Computer Science* 5(2):1–38, 2009.

[11] Hansson, Sven Ove, *A Textbook of Belief Dynamics. Theory Change and Database Updating*, Kluwer, 1999.

[12] Kratzer, Angelika, 'What *must* and *can* must and can mean', *Linguistics and Philosophy* 1:337–355, 1977.

[13] Leitgeb, Hannes, and Krister Segerberg, 'Dynamic doxastic logic: why, how and where to?', *Synthese* 155(2):167–190, 2007.

[14] Lewis, David, *Counterfactuals*, Blackwell Publishers, Oxford, 1973.

[15] Liu, Fenrong, *Reasoning about Preference Dynamics*, vol. 354 of *Synthese Library*, Springer-Verlag, 2008.

[16] Liu, Fenrong, 'A two-level perspective on preference', *Journal of Philosophical Logic*, 2011. To Appear.

[17] Moss, Larry, and Rohit Parikh, 'Topological reasoning and the logic of knowledge', in Yoram Moses, (ed.), *Proceedings of TARK IV*, Morgan Kaufmann, 1992.

[18] Pacuit, Eric, 'Neighborhood semantics for modal logic: An introduction', 2007. ESSLLI 2007 course notes (`ai.stanford.edu/~epacuit/classes/`).

[19] Pauly, Marc, *Logic for Social Software*, Ph.D. thesis, ILLC University of Amsterdam Dissertation Series DS 2001-10, 2001.

[20] Plaza, J., 'Logics of public communications', *Synthese: Knowledge, Rationality, and Action* 158(2):165–179, 2007.

[21] Rott, Hans, *Change, Choice and Inference: A Study in Belief Revision and Nonmonotonic Reasoning*, Oxford University Press, 2001.

[22] Rott, Hans, 'Shifting priorities: Simple representations for 27 iterated theory change operators', in H. Lagerlund, S. Lindström, and R. Sliwinski, (eds.), *Modality Matters: Twenty-Five Essays in Honor of Krister Segerberg*, vol. 53 of *Uppsala Philosophical Studies*, 2006, pp. 359–384.

[23] Segerberg, Krister, 'Belief revision from the point of view of doxastic logic', *Journal of the IGPL* 3(4):535–553, 1995.

[24] Shafer, Glen, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.

[25] Shoham, Yoav, and Kevin Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, 2009.

[26] Stalnaker, Robert, 'Knowledge, belief and counterfactual reasoning in games', *Economics and Philosophy* 12(02):133–163, 1996.

[27] Su, Kaile, Abdul Sattar, Guido Governatori, and Qingliang Chen, 'A computationally grounded logic of knowledge, belief and certainty', in *Proceedings of the fourth international joint conference on autonomous agents and multiagent systems*, AAMAS '05, 2005, pp. 149–156.

[28] van Benthem, Johan, 'Dynamic logic for belief revision', *Journal of Applied Non-Classical Logics*, 14(2):129–155, 2004.

[29] van Benthem, Johan, 'A note on modeling theories', in R. Festa, A. Aliseda, and J. Peijnenburg, (eds.), *Poznan Studies in the Philosophy of the Sciences and Humanities: Confirmation, Empirical Progress and Truth Approximation. Essays in Debate with Theo Kuipers*, vol. 17, 2005, pp. 403–419.

[30] van Benthem, Johan, 'Merging observation and access in dynamic logic', *Studies in Logic* 1(1):1–17, 2008.

[31] van Benthem, Johan, *Logical Dynamics of Information Flow*, Cambridge University Press, 2011.

[32] van Benthem, Johan, and Stefan Minica, 'Toward a dynamic logic of questions', in Xiangdong He, John F. Horty, and Eric Pacuit, (eds.), *Logic, Rationality, and Interaction*, vol. 5834 of *LNCS*, Springer, 2009, pp. 27–41.

[33] van Benthem, Johan, and Eric Pacuit, 'Dynamic logics of evidence-based belief', Tech. rep., Institute for Logic, Language and Computation, 2011.

[34] van Benthem, Johan, and Fernando Raymundo Velázquez-Quesada, 'The dynamics of awareness', *Synthese* 177:5–27, 2010.

[35] van Ditmarsch, H., W. van der Hoek, and B. Kooi, *Dynamic Epistemic Logic*, Synthese Library, Springer, 2007.

[36] Velazquez-Quesada, F. R., 'Inference and update', *Synthese (Knowledge, Rationality and Action)* 169(2):283–300, 2009.

[37] Veltman, Frank, 'Prejudices, presuppositions and the theory of conditionals', in Jeroen Groenendijk, and Martin Stokhof, (eds.), *Amsterdam Papers in Formal Grammar*, vol. 1, 1976, pp. 248–281.

[38] Zvesper, Jonathan, *Playing with Information*, Ph.D. thesis, ILLC University of Amsterdam Dissertation Series DS-2010-02, 2010.

Johan van Benthem
Institute for Logic, Language and Computation
University of Amsterdam, 1090 GE
Amsterdam, The Netherlands
Department of Philosophy
Stanford University
Stanford, CA 94305, USA
J.vanBenthem@uva.nl

Eric Pacuit
Tilburg Center for Logic and Philosophy of Science
Tilburg University
P.O. Box 90153, 5000 LE
Tilburg, The Netherlands
e.j.pacuit@uvt.nl